Human Factors in Robots, Drones and Unmanned Systems, Vol. 93, 2023, 1–8
https://doi.org/10.54941/ahfe1003741

Navigating Team Cognition: Goal Terrain as Living Map to Situation Awareness

Melissa J. Carraway[1,2], Victoria Chang[1,2], Nathan P. Hahn[3], and Susan G. Campbell[1,2]

[1]Applied Research Laboratory for Intelligence and Security, University of Maryland, College Park, MD 20740, USA
[2]College of Information Studies, University of Maryland, College Park, MD 20742, USA
[3]DEVCOM Army Research Laboratory, Aberdeen Proving Ground, MD, USA

ABSTRACT

Effective human-machine communication relies on shared goals and relaying of intent between human and machine agents. Based on insights from human-machine teaming and the human teamwork literature, we propose a naturalistic approach to communication that relies on shared representation of goals. This concept, which we call Goal Terrain, provides both a means of supporting situation awareness and a method for communicating and highlighting updates to a predetermined plan. In this paper, we lay out some requirements for Goal Terrain and propose ways that it can be used to organize communication within human-agent teams.

Keywords: Human-autonomy team, Autonomy, Team, Human factors, Robotics, User interface, Interaction design

INTRODUCTION

Human-autonomy teams will require new methods of communication and coordination that reinvent interaction between humans and agents from an operator model, where humans explicitly control a system and only permit limited or supervised autonomy, to a teammate model, where agents can make certain decisions. This should allow human teammates to be concerned with the work of the team rather than supervising the agents in it (de Visser et al., 2018). Deficits remain that preclude us from fully realizing effective teaming with autonomous systems, however, especially for “hybrid” teams that include both humans and machine agents. One major impediment to agents assisting human users lies in the mutually intelligible communication of intent. How can we design agents to function in a more naturalistic way, with the ability to reason better about shared goals, and to perform necessary teamwork skills?

Here, we define teaming as two or more individuals who perform organizationally relevant tasks, share common goals, have task interdependence, and maintain and manage boundaries and roles (Kozlowski and Bell, 2003). Human-autonomy teams fulfill these requirements with a mix of human and autonomous agent teammates. Other roles for autonomous systems exist, such as supporting individual human teammates or supporting specific team functions (Sycara & Sukthankkar, 2006), but this paper focuses on autonomous teammates.

In contrast to teleoperated systems and other tools, autonomous teammates have the potential to respond in off-nominal situations, to solve problems, and to adapt and act in a proactive fashion as essential for teamwork (Lyons at al., 2021). Lyons and colleagues (2021) also highlight coordination, cooperation, and communication as key team processes that enable human-autonomy teaming.

Some of the coordination that occurs in human-human teams is tacit and relies on shared metacognitive activities and shared mental models; systems enabling human-autonomy teams would need to make that coordination explicit to enable adaptation and directability for non-human team members (Sycara & Sukthankar, 2006). Shared understanding of goals and the current environment also enable team adaptation (Salas, Sims, & Burke, 2005).

Another requirement for human-autonomy teams is bi-directional communication, where autonomous teammates provide information to human team members and human team members provide information and guidance to autonomous teammates. This communication requires a shared representation of goals, intent, and expected outcomes (Barnes, et al., 2021). Ideally, such communication would not require the human to learn new ways of communicating but could leverage naturalistic communication styles.

One framework for choosing the information that autonomous teammates should share with the rest of a team to develop bidirectional communication comes from Chen's Situation Awareness-based Transparency (SAT) model, which specifies which information an agent should share to support different levels of human SA about a single agent's status, reasoning process, and projections about future consequences of the current situation (Chen, 2014).

Effective team communication requires (and creates) shared mental models and shared SA, and these components are especially important for human-autonomy teams (Ososky, et al., 2012). These shared mental models include models of the task, of the environment, or of the capabilities of the team (both individually and collectively). Shared understanding of team goals and how individual tasks fit into them is key to team performance (Salas, Sims, & Burke, 2005).

Situation awareness can be conceptualized with three different levels: the first is perception of elements in the environment, the second is awareness of the impacts of those elements, and the third is projecting the impacts in the future (Endsley, 1988). Some controversy exists around whether SA is a term that applies to autonomous systems, with some authors arguing that autonomous teammates exist to reduce the burden of taskwork on human teammates, leaving those humans with the primary task of maintaining SA (Ososky, et al., 2012). Regardless of whether SA is unique to the human members of a team, the relevant components must be made visible to human team members, and shared mental models are required to coordinate translation between the levels of SA.

AI teammates should function to support teamwork, focusing on leveraging and supporting team awareness and driving shared SA. Teaming can be made more effective and efficient when teammates are aware of the situation as it unfolds and can be responsive to changes in the environment, objectives, or task constraints. Each level of SA is critical at both the individual and team levels, and can be supported appropriately by a graphic user interface (GUI). First, the first level of SA can be supported with appropriately coded alerts, alarms and notifications that bring attention to relevant and dynamic issues of the environment. Second, the AI can be designed to contextualize data to form information that is relevant specifically to the user who receives it. Finally, AI can project future outcomes that allow operators to make decisions based on the impact that events and changes in the environment will have on those processes in motion and how that impacts their goals. Instantiating the requirements for a human-autonomy team that will accomplish these goals is the main thrust of goal terrain.

Goal Terrain

Our logical first step for developing communication methods that may elevate AI to the level of a teammate is to support aspects of teamwork that foster team awareness. Challenges include the development of team-based affordances for the development of shared awareness (Lyons, 2021), displays that illustrate a common picture of the situation that enable users to understand how it is evolving and then to evaluate solutions to problems as they arise (Roth et al., 2018), and to enable bi-directional communication between human and machine agents (Wright, 2022). This is especially challenging given that the rate at which the environment and situations change creates a need for information beyond that which a human can reasonably process (Endsley & Jones 2012). How can we communicate all of this in a way that is natural, intuitive, and elegant?

Because so much of what needs to be understood and conveyed amounts to SA, we begin with the situation awareness-based agent transparency (SAT) model developed by Chen et al., (2014). This model describes the information that should be displayed by a user interface to support each level of SA. The SAT model decomposes level one SA into goals and actions, the agent's current status and plans, and supporting information such as purpose, process, performance, and perception. Level two SA is described as the agent's reasoning, and is broken down into processes, motivations, and constraints in the environment. Finally, the third level of SA, which consists of projections, includes agent predictions and any uncertainty ascribed to them. While the SAT model frames the prerequisites for conveying SA, it stops short of describing how best to highlight changes in the environment or to use it to foster negotiations and discussion. To this end, this paper expands on that concept to develop the concept of goal terrain as a visual representation of a plan and its projected outcome that can be used to communicate bi-directionally between humans and AI. See *Table 1* for details.

Plan Articulation. The concept of goal terrain (GT) expands on the SAT model to develop the shared visualization to support team cognition, affordances from which negotiation can stem, highlight changes to the situation or environment, and project future outcomes. Because information should be

Table 1. Key communication mappings for goal terrain.

Key Communications	Visualization Requirements
Plan articulation	• Goals, subgoals • Dependencies between goals • Resources (material or human) • Constraints (such as time) Expectations (such as weather forecast)
Changes in the environment or situation	• Deltas between expectation and reality Deviations from the plan in terms of time, resources, enemy forces, etc.
Impacts on future states and goal achievement	• Updated probability of goal achievement • Indication of relevant changes in the environment or constraints that negatively impacts goal completion • Areas of uncertainty due to lack of information Risk assessment

organized in terms of an individual's major goals (Endsley & Jones, 2012), goal terrain begins with an articulation of a plan that is framed by those goals that are relevant to the viewer. Each goal is broken down and represented by its component sub-goals. Details about these goals include constraints, resources allocated to them, dependencies between goals, and their relative priorities. Plan articulations also include those aspects that impact feasibility but include relevant expectations, such as weather forecasts or the absence of enemy forces.

Changes in the Environment or Situation. It is said that no plan survives first contact, and that the enemy always has a say in what transpires. There is always a difference between what is expected to happen and what reality brings. For that reason, an important aspect of goal terrain is how the AI is expected to track deltas between what was planned and what happens. These deltas can manifest in changes to resources, missed dependencies, or changes in the environment. The difference between what was expected and what has happened will be displayed as important information to the relevant users and will be especially important as these deltas impact the projection of future states.

Projection of Future States. Endsley's (1988) concept of third-level SA involves the ability to not only recognize that something has happened, but to understand its importance now and to use that information to make projections into the future. Different projections can stem from variations in understanding of available resources, constraints, or capabilities. For this reason, simply representing a plan and its constraints are not enough: it is the relationship between the goals, constraints, and relevant events that create an accurate projection, and fosters not only a scaffold for the development of shared mental models, but also a frame from which to make informed decisions for the purposes of adaptation and backup behaviors. Goal terrain can be used as a visual and logical artifact for the purposes of understanding tradeoffs between courses of action (COAs) as plans change in the dynamic environment.

The ability to project future states is also critical for the identification of risk and the quantification of uncertainty. The topic of risk includes many things, including resource availability, the brittleness of a plan, the ability to meet time constraints, and others. Areas of uncertainty also contribute to the larger picture of risk and are often much harder to describe and quantify.

COA Comparisons. It is natural for humans to discuss plans by way of projected outcomes, especially when negotiating between options. The organizing principle of Goal Terrain is the idea that the goals of a team can be represented in abstract space, along with multiple options to completion, as "better" paths might be highlighted. Since machines will eventually be capable of weighing alternatives and making decisions based on both internal and external factors (Ososky et al., 2012), it is only natural that an AI would have the ability to automatically project and grade actions based on projected outcomes. This may prove to be more natural as a means of human-to-ai discourse, as people naturally engage in satisficing by way of heuristics to come up with the first 'good enough' option instead of exhausting all potential alternatives (Simon, 1955). In the case of a goal that cannot be completed as planned, the GUI might offer suggestions on a different resource or method to achieve the goal and suggest them to the user, framing the differences between options first in the projected outcomes of making the change and then by showing the details of it.

Alternatively, team members may be able to change or lessen constraints on a goal to create more options. Both options present value as methods for a user to make sense of their current plan and environment and situation while also understanding the play space. This should support the development of shared mental models by way of the explorative negotiation process and seeing how the AI "projects" its understanding of the future. It also illustrates parameters, risks, uncertainties, and areas of vulnerability.

Information Overload. Communication does not always happen as intended, either because it was not received or because it was not interpreted correctly. This can happen as a result of too much data (information overload), or because the information is not conveyed at the right time for it to be integrated (Salas, Sims, & Burke, 2005). This problem can be alleviated either by providing information either via data dashboards or by providing teammates with information that can be pushed to relevant members as needed (Demir et al., 2017). Because what is salient and important changes on the context of the use case, the team, and the individual, graphical user interfaces of dense information must be properly curated for different users and user types to reduce the cognitive load of sifting through and translating data into something that is usable and actionable. It is our hope that goal terrain can help to make the right information salient at the right time for the right individuals to quickly make informed decisions.

A Goal Terrain Example

In preparation for an approaching convoy of vehicles, a scout search team conducts a search for a reported suspicious package on a nearby busy roadway. The scout team consists of one human member, two ground robots, and one UAV. Dependencies include that the roadway must be cleared before the convoy can proceed. Constraints include that the human team member must remain a safe distance from the package considering the potentially dangerous nature of the item, and that the search must be completed before the convoy arrives. Competing goals include potential simultaneous tasking. For example, the UAV can either monitor for enemy combatants or search for a remote detonator, but not both. The goals of conducting a sweep, identifying and potentially neutralizing the suspicious package, and expectations of constraints and time requirements are encoded into the model, and will be monitored by the Goal Terrain Agent. Also encoded are the makeup of the team, their relevant responsibilities, statuses, and their own constraints. Finally, certain relevant expectations, to include the planned arrival time of the convoy, available tools, and number of reported packages that are expected to be found will also be captured.

The goal terrain visualization is meant to be a representation of the evolving goal space, its critical aspects include the deltas that occur between what was planned and how recent events change projections on plan feasibility, risk, and timing. It also proffers multiple alternatives of how to "fix" the plan, as well as tradeoffs between those alternatives. In our example, the convoy has a goal to reach its destination safely, and by a certain time. First, and most importantly, the discovery of a suspicious package creates new risk, not only to the forces who may encounter it, but logistically with regards to the goal of the convoy's safe and timely travel. While there is a chance that the suspicious package is not a threat, and that travel may not be impacted, it is still possible that this package may take time to disable, or that if it detonates, may make the road unpassable. Second, discovery of the package creates new goals that include evaluation and disposal of the package. Finally, the resources of the team are somewhat limited. For instance, if a UAV is scouting

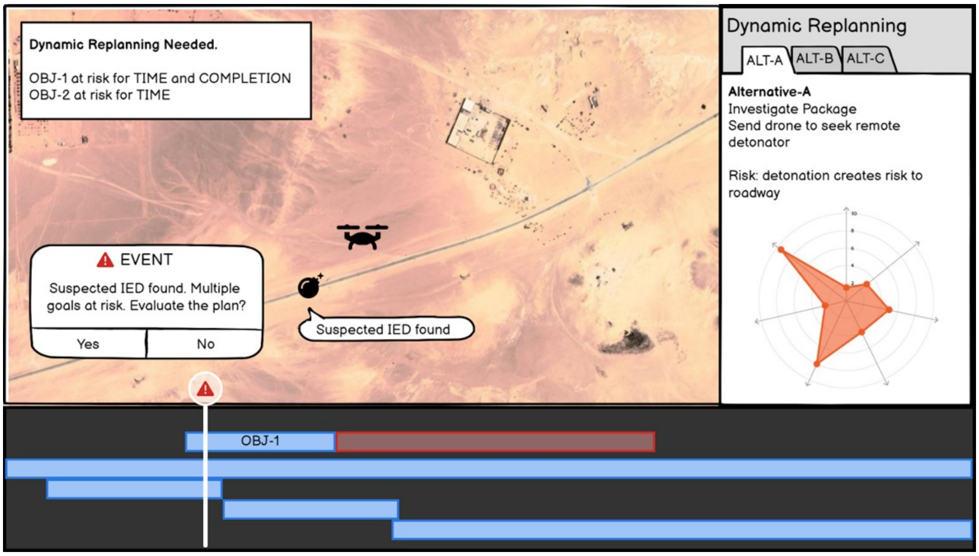

Figure 1: Rough concept drawing visualizing goal terrain.

for a remote detonator, they are unable to do other tasking which might be to keep watch for ambush. Working with limited resources means that choices must be made, and tradeoffs accepted.

As the task continues towards completion, the actual values of the situation are monitored and compared against what was encoded by the system. In the event of a material change to the plan, for example if multiple suspicious packages are found instead of one, the system will provide feedback, alternatives, and notifications to allow the team to make adaptations as necessary to reach the objectives. In this case, alternatives may include the recruitment of more resources for neutralizing suspicious packages, slowing or diverting the convoy. Based on the needs of the mission, the team can quickly make judgments about the task and reach a satisfactory solution.

Research Gaps

This paper represents the first step into the exploration of goal terrain by way of displaying the information of a plan as centered on goals and projections as to their outcomes and likelihood of success. It illustrates dynamics in the environment in such a way as to not only highlight the deltas between what was expected and what is unfolding, but to frame those changes in terms of those goal projections. It is expected that this will increase human and machine comprehension of a plan and a plan's options.

Future work would do well to articulate and test methods for visualizing these concepts. Creation of an interaction method and GUI for a specific use case would allow for the testing of theories involving the development of shared mental models and decision making abilities as supported by such a tool. Finally, future work could test whether these methods positively impact replanning ability, adaptability, and resilience in team operations when working in uncertain environments.

CONCLUSION

Using Goal Terrain, machine and human agents will be able to mutually support one another, create a greater sense of mutual predictability, adapt to dynamic needs, and negotiate next steps. Human and machine agents alike will be able to proactively offer mutual assistance to one another by better monitoring individual and team activity and the relationships between events, resources, and goals. Agents will have a greater ability to offer insight and ask for guidance as appropriate, as they will have more information to work from regarding plan execution. Providing autonomous agents with enhanced context and a shared representation of goals and progress that they can query will reduce human interaction and intervention requirements while ensuring that those questions that are asked are more impactful and informative. Agents with the ability to manipulate and represent goals in common with human agents should help to mitigate the problems of mixed-initiative interactions by managing the uncertainties of agent goals, focus of attention, plans, and status. In short, we hope this might improve the speed and efficiency of teammates as they adapt to a changing environment, increase visibility into and accounting for critical tradeoffs between COAs, and offer insights into resource management that may not be immediately apparent.

ACKNOWLEDGMENT

This work was sponsored under U.S. Army Cooperative Agreement W911NF2120076.

REFERENCES

Barnes, M. J., Wang, N., Pynadath, D. V. and Chen, J. Y., 2021. Human-agent bidirectional transparency. *Trust in Human-Robot Interaction* (pp. 209–232). Academic Press.

Chen, J. Y. C., Procci, K., Boyce, M., Wright, J., Garcia, A., & Barnes, M. J. 2014. Situation awareness-based agent transparency (Tech Rep. ARL-TR-6905). *Aberdeen Proving Ground, MD: U. S. Army Research Laboratory.*

de Visser, E. J., Pak, R. and Shaw, T. H., 2018. From 'automation' to 'autonomy': the importance of trust repair in human–machine interaction. Ergonomics, 61(10), pp. 1409–1427.

Demir, M., McNeese, N. J. and Cooke, N. J., 2017. Team situation awareness within the context of human-autonomy teaming. *Cognitive Systems Research, 46*, pp. 3–12.

Endsley, M. R., 1988, October. Design and evaluation for situation awareness enhancement. *In Proceedings of the Human Factors Society annual meeting* (Vol. 32, No. 2, pp. 97–101). Sage CA: Los Angeles, CA: Sage Publications.

Endsley, M. R., Jones, D. G., 2012. Designing for situation awareness: An approach to user-centered design (2nd ed.). CRC press.

Kozlowski, S. W. J., BELL (2003), '*Work groups and teams in organizations'. Handbook of Psychology: Industrial and Organizational Psychology.* Wiley, London, 12, pp. 333–375.

Lyons, J. B., Sycara, K., Lewis, M., & Capiola, A. 2021. Human–Autonomy Teaming: Definitions, Debates, and Directions. *Frontiers in Psychology*, *12*(May), 1–15. https://doi.org/10.3389/fpsyg.2021.589585

Ososky, S., Schuster, D., Jentsch, F., Fiore, S., Shumaker, R., Lebiere, C., Kurup, U., Oh, J., & Stentz, A. 2012. The importance of shared mental models and shared situation awareness for transforming robots from tools to teammates. *Unmanned Systems Technology XIV*, 8387, 838710. https://doi.org/10.1117/12.923283

Roth, E., Depass, B., Harter, J., Scott, R. and Wampler, J., 2018, September. Beyond levels of automation: Developing more detailed guidance for human automation interaction design. *In Proceedings of the Human Factors and Ergonomics Society Annual Meeting* (Vol. 62, No. 1, pp. 150–154). Sage CA: Los Angeles, CA: SAGE Publications.

Salas, E., Sims, D. E., & Shawn Burke, C. 2005. Is there a "Big Five" in Teamwork? *Small Group Research, 36*(5), pp. 555–599.

Salas, E., Sims, D. E., & Shawn Burke, C. 2005. Is there A "big five" in teamwork? *Small Group Research, 36*(5), *555–599.*

Simon, H. A., 1955. A behavioral model of rational choice. The quarterly journal of economics, pp. 99–118.

Sycara, K. and Sukthankar, G., 2006. Literature review of teamwork models. *Robotics Institute, Carnegie Mellon University, 31*(31), pp. 1–31.

Wright, J. L., Lakhmani, S. G., Schwartz, M. R. and Chen, J. Y., 2022. B*idirectional Communications in Human-Agent Teaming: The Effect of Communication Style.* DEVCOM Army Research Laboratory.